

Self-Reflective Risk-Aware Artificial Cognitive Modeling for Robot Response to Human Behaviors

Fei Han^a, Christopher Reardon^b, Lynne E. Parker^b, Hao Zhang^a

^a*Human-Centered Robotics Lab, Department of Electrical Engineering and Computer Science,
Colorado School of Mines, Golden, CO 80401, USA*

^b*Distributed Intelligence Lab, Department of Electrical Engineering and Computer Science,
University of Tennessee, Knoxville, TN 37996, USA*

Abstract

In order for cooperative robots (“co-robots”) to respond to human behaviors accurately and efficiently in human-robot collaboration, interpretation of human actions, awareness of new situations, and appropriate decision making are all crucial abilities for co-robots. For this purpose, the human behaviors should be interpreted by co-robots in the same manner as human peers. To address this issue, a novel interpretability indicator is introduced so that robot actions are appropriate to the current human behaviors. In addition, the complete consideration of all potential situations of a robot’s environment is nearly impossible in real-world applications, making it difficult for the co-robot to act appropriately and safely in new scenarios. This is true even when the pretrained model is highly accurate in a known situation. For effective and safe teaming with humans, we introduce a new generalizability indicator that allows a co-robot to self-reflect and reason about when an observation falls outside the co-robot’s learned model. Based on topic modeling and two novel indicators, we propose a new *Self-reflective Risk-aware Artificial Cognitive* (SRAC) model. The co-robots are able to consider action risks and identify new situations so that better decisions can be made. Experiments both using real-world datasets and on physical robots suggest that our SRAC model significantly outperforms the traditional methodology and enables better decision making in response to human activities.

Email addresses: fhan@mines.edu (Fei Han), creardon@utk.edu (Christopher Reardon), leparker@utk.edu (Lynne E. Parker), hzhang@mines.edu (Hao Zhang)

Keywords:

Reflective cognitive models, risk-aware decision making, human behavior interpretation, human-robot collaboration

1. Introduction

In human-robot collaboration, it is crucial for a cooperative robot (“co-robot”) to have the abilities of perception of human activities and corresponding appropriate decision-making to understand and interact with human peers. In order to provide these important capabilities, an artificial cognitive model integrating perception, reasoning, and decision making modules is required by intelligent co-robots to respond to humans effectively. Artificial cognition has its origin in cybernetics; its intention is to create a science of mind based on logic [1]. Among other mechanisms, cognitivism is a most widely used cognitive paradigm [2]. Several cognitive architectures were developed within this paradigm, including ACT-R [3] (and its extensions ACT-R/E [4], ACT-R Φ [5], etc), Soar [6], C4 [7], and architectures for robotics [8]. Because an architecture represents the connection and interaction of different cognitive components, it cannot accomplish a specific task on its own without specifying each component that can provide knowledge to the cognitive architecture. The combination of the cognitive architecture and components is usually referred to as a cognitive model [2].

Implementing such an artificial cognitive system is challenging, since the high-level processes (e.g., reasoning and decision making) must be able to seamlessly work with the low-level components, e.g., perception, under significant uncertainty in a complex environment [9]. In the context of human-robot collaboration, perceiving human behaviors is a necessary component, where uncertainty arises due to human behavior complexity, including variations in human motions and appearances, and challenges of machine vision, such as lighting changes and occlusion. This perception uncertainty is addressed in this work using the bag-of-visual-words (BoW) representation based on local spatio-temporal features, which has previously shown promising performance [10, 11, 12].

To further process the perceptual data, a high-level reasoning component is nec-

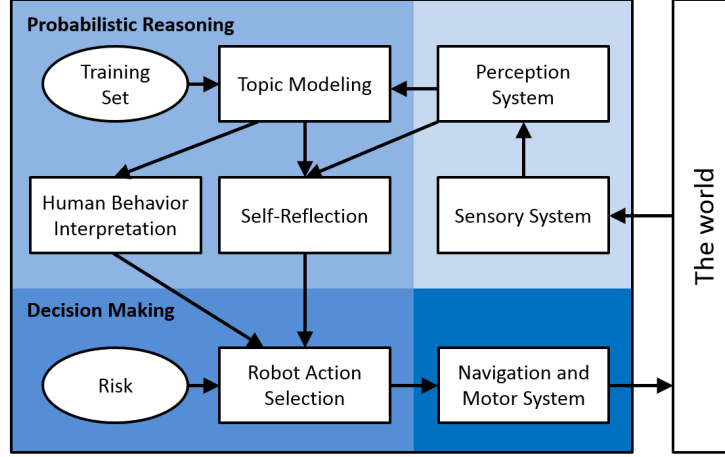


Figure 1: Overview of the SRAC model for robot response to human activities. The novel self-reflection module allows a co-robot to reason about when the learned knowledge no longer applies. Decisions are made by considering both human activity category distributions and robot action risks. Entities in ellipses are prior knowledge to the SRAC model. Information flows from modules with lighter colors to those with darker colors.

essary for a co-robot to make decisions. In recent years, topic modeling has attracted increasing attention in human behavior discovery and recognition due to its ability to generate a distribution over activities of interest, and its promising performance using BoW representations in robotics applications [12, 13]. However, previous work only aimed at human behavior understanding; the essential task of incorporating topic modeling into cognitive decision making (e.g., selecting a response action) is not well analyzed.

Traditional activity recognition systems typically use accuracy as a performance metric [14]. Because the accuracy metric ignores the distribution of activity categories, which is richer and more informative than a single label, it is not appropriate for decision making. For example, in a task of behavior understanding with two categories, assume that two recognition systems obtain two distributions $[0.8, 0.2]$ and $[0.55, 0.45]$ on a given observation, and the ground truth indicates the first category is correct. Although both systems are accurate, in the sense that the most probable category matches the ground truth, the first model obviously performs better, since it better separates the

correct from the incorrect assignment. Previous studies did not consider this important phenomenon.

In real-world applications, artificial cognitive models must be applied in an online fashion. If a co-robot is unable to determine whether its knowledge is accurate, then if it observes a new human behavior that was not presented during the training phase, it cannot be correctly recognized, because the learned behavior recognition model no longer applies. Decision making based on incorrect recognition in situations like these can result in inappropriate or even unsafe robot action response. Thus, an artificial cognitive model requires the capability to self-reflect whether the learned activity recognition system becomes less applicable, analogous to human self-reflection on learned knowledge, when applied in a new unstructured environment. This problem was not well investigated in previous works.

In this paper, we develop a novel artificial cognitive model, based on topic models, for robot decision making in response to human behaviors. Our model is able to incorporate human behavior distributions and take into account robot action risks to make more appropriate decisions (we label this “risk-aware”). Also, our model is able to identify new scenarios when the learned recognition subsystem is less applicable (we label this “self-reflective”). Accordingly, we call our model the *self-reflective, risk-aware artificial cognitive* (SRAC) model.

Our primary contributions are twofold:

- Two novel indicators are proposed. The *interpretability indicator* (I_I) enables a co-robot to interpret category distributions in a similar manner to humans. The online *generalizability indicator* (I_G) measures the human behavior recognition model’s generalization capacity (i.e., how well unseen observations can be represented by the learned model).
- A novel artificial cognitive model (i.e., SRAC) is introduced based on topic models and the indicators, which is able to consider robot action risks and perform self-reflection to improve robot decision making in response to human activities in new situations.

The rest of the paper is organized as follows. We first overview the related work

in Section 2. Then, the artificial cognitive architecture and its functional modules are described in Section 3. Section 4 introduces the new indicators. Section 5 presents self-reflective risk-aware decision making. Experimental results are discussed in Section 6. Finally, we conclude our paper in Section 7.

2. Related Work

In this section, we provide an overview of a variety of methods related to our proposed SRAC cognitive model for human-robot teaming, including human activity recognition, topic models, and artificial cognitive modeling.

2.1. Human Activity Recognition

We focus our review on the commonly used sequential and space-time volume methods [15]. A comprehensive review of different aspects of human activity recognition (HAR) is presented in [15] and [16].

A popular sequential method is to use centroid trajectories to identify human activities in visual data, in which a human is represented as a single point indicating the human’s location. Chen and Yang represented a human with just a point to derive the gait features for the pedestrian detection [17]. Ge *et al.* extracted pedestrian trajectories from video to automatically detect small groups of people traveling together [18]. These methods can avoid the influence of human appearances such as dresses and carrying, but are not able to recognize activities involving various relative body-part movements. Another sequential method relies on human shapes, including human contours and silhouettes. Singh *et al.* extracted directional vectors from the silhouette contours and utilize the distinct data distribution of these vectors in a vector space for activity recognition [19]. Junejo *et al.* transformed silhouettes of a human from every frame into time-series, then each of these time series is converted into the symbolic vector to represent actions [20]. A third method is based on body-part models. Zhang and Parker implemented a bio-inspired predictive orientation decomposition (BIPOD) to construct representations of people from skeleton trajectories for the activity recognition and prediction, where the human body is decomposed into five body parts [21].

However, techniques based on shapes and/or body-part models rely on human and body-part detection, which are hard-to-solve problems due to occlusions and dynamic backgrounds, among others.

Space-time volume methods use local features to represent local texture and motion variations regardless of global human appearance and activity. A large number of HAR methods are based on SIFT features [22] and its extensions [23]. For example, Behera *et al.* proposed a random forest that unifies randomization, discriminative relationships mining and a Markov temporal structure for real-time activity recognition with SIFT features [24]. However, SIFT features only encode appearance information and are not able to represent temporal information. STIP features were introduced in [25] and SVMs were applied to classify human activities. Dollar *et al.* used separable filters in spatial and temporal dimensions to extract features for HAR [26]. Four-dimensional features were also introduced in [12] to combine depth information to classify human activities.

Previous work focused only on recognizing human activities but did not discuss the consequent issue: how a co-robot can make decisions based on recognition results, especially when risks are associated with different human activities.

2.2. Topic Models and Evaluation

Among other machine learning techniques, topic modeling has been widely applied to HAR. Zhao *et al.* incorporated Bayesian learning into an undirected topic model and proposed a "relevance topic model" for the unstructured social group activity recognition [27]. A semi-latent topic model trained in a supervised fashion was introduced in [28] and used to classify activities in videos. Zhang and Parker adopted topic models to classify activities in 3D point clouds from color-depth cameras on mobile robots [12]. Topic models were also widely used to discover human activities in streaming data. The use of topic models was explored in [29] to discover daily activity patterns in wearable sensor data. An unsupervised topic model was proposed in [30] to detect daily routines from streaming location and proximity data. Taking temporal and/or object relational information into account, Freedman *et al.* explored a new method using topic models for both plan recognition and activity recognition objective [31].

Although there is a significant body of work introducing and developing sophisticated topic models, few efforts have been undertaken to evaluate them. Existing methods are dominated by either intrinsic methods, (e.g., computing the probability of held-out documents to evaluate generalization ability [32]) or extrinsic methods that rely on external tasks, (e.g., information retrieval [33]). Some work also focused on evaluation of topic modeling’s interpretability as semantically coherent concepts. For example, Chang *et al.* demonstrated that the probability of held-out documents is not always a good indicator of human judgment [34]. Newman *et al.* showed that metrics based on word co-occurrence statistics are able to predict human evaluations of topic quality [35].

As recently pointed out by Blei [36], topic model evaluation is an essential research topic. Despite this, previous works use only the accuracy metric to evaluate topic modeling results in HAR tasks; issues such as the model’s interpretability and generalizability have not been studied. In this paper, we analyze these two aspects of topic model evaluation in HAR tasks, explore their relationship, and show how they can be used to improve robot decision making.

2.3. Artificial Cognitive Modeling

Artificial cognition has its origin in cybernetics with the intention to create a science of mind based on logic [1]. Among other cognitive paradigms, cognitivism has undoubtedly been predominant to date [2]. Within the cognitivism paradigm, several cognitive architectures were developed, including Soar [6], ACT-R [3] (and its extensions ACT-R/E [4], ACT-R Φ [5], etc), C4 [7], and architectures for robotics [8], which are relatively independent of applications [37]. Because architectures represent the mechanism for cognition but lack the relevant information for using that mechanism, they cannot accomplish anything in their own right and need to be provided with knowledge to conduct a specific task. The combination of a cognitive architecture and a particular knowledge set is generally referred to as a *cognitive model* [2]. The knowledge incorporated in cognitive models is typically determined by human designers [2]. The knowledge can be also learned and adapted using machine learning techniques.

Cognitive models have been widely used in human-machine interaction and robotic

vision applications. For example, cognitive modeling was adopted in [38, 39, 40] to construct intelligent human-machine interaction systems. Cognitive perception systems were also used to recognize traffic signs [41, 42], interpret traffic behaviors [43, 44], and recognize human activities [45, 46]. Over the last decade, probabilistic models of cognition, as an alternative of deterministic cognitive models, have attracted more attention in cognitive development [47]. For example, an adaptive remote data mirroring system was proposed applying dynamic decision networks in [48]. Another cognitive model was introduced in [49] to apply dynamic Bayesian networks for vehicle classification. Probabilistic models have also been widely used for learning and reasoning in cognitive modeling [50].

We believe we are the first to adopt topic models for the construction of reliable artificial cognitive models and show that they are particularly suited for this task. We demonstrate topic modeling’s ability to combine risks in decision making. In addition, we develop two evaluation metrics and show their effectiveness in model selection and decision making. These aspects were not addressed in previous artificial cognitive modeling research.

3. Topic Modeling for Artificial Cognition

3.1. Cognitive Architecture Overview

The proposed SRAC model is inspired by the C4 cognitive architecture [7]. As shown in Fig. 1, our model is organized into four modules by their functionality:

- *Sensory and perception*: Visual cameras observe the environment. Then, the perception system builds a BoW representation from raw data, which can be processed by topic models.
- *Probabilistic reasoning*: Topic models are applied to reason about human activities, which are trained off-line and used online. The training set is provided as a prior that encodes a history of sensory information. This module uses the proposed indicators to select topic models that better match human’s perspective, and to discover new activities in an online fashion.

- *Decision making*: Robot action risk based on topic models and the evaluation results is estimated and a response robot action that minimizes this risk is selected. The risk is provided as a prior to the module.
- *Navigation and motor system*: The selected robot action is executed in response to human activities.

3.2. Topic Modeling

Latent Dirichlet Allocation (LDA) [51], which showed promising activity recognition performance in our prior work [12], is applied in the SRAC model.

Given a set of observations $\{\mathbf{w}\}$, LDA models each of K activities as a multinomial distribution of all possible visual words in the dictionary \mathbf{D} . This distribution is parameterized by $\boldsymbol{\varphi} = \{\varphi_{w_1}, \dots, \varphi_{w_{|\mathbf{D}|}}\}$, where φ_w is the probability that the word w is generated by the activity. LDA also represents each \mathbf{w} as a collection of visual words, and assumes that each word $w \in \mathbf{w}$ is associated with a latent activity assignment z . By applying the visual words to connect observations and activities, LDA models \mathbf{w} as a multinomial distribution over the activities, which is parameterized by $\boldsymbol{\theta} = \{\theta_{z_1}, \dots, \theta_{z_K}\}$, where θ_z is the probability that \mathbf{w} is generated by the activity z . LDA is a Bayesian model, which places Dirichlet priors on the multinomial parameters: $\boldsymbol{\varphi} \sim \text{Dir}(\boldsymbol{\beta})$ and $\boldsymbol{\theta} \sim \text{Dir}(\boldsymbol{\alpha})$, where $\boldsymbol{\beta} = \{\beta_{w_1}, \dots, \beta_{w_{|\mathbf{D}|}}\}$ and $\boldsymbol{\alpha} = \{\alpha_{z_1}, \dots, \alpha_{z_K}\}$ are the concentration hyperparameters.

One of the major objectives in HAR tasks is to estimate the parameter $\boldsymbol{\theta}$, i.e., the per-observation activity proportion. However, exact parameter estimation is intractable in general [51]. Our model applies Gibbs sampling [52] to compute the *per-observation activity distribution* $\boldsymbol{\theta}$, based on two considerations: 1) This sampling-based method is generally accurate, since it asymptotically approaches the correct distribution [53], and 2) This method can be used to intrinsically evaluate topic model's performance [32], thereby providing a consistent method to infer, learn, and evaluate topic models. At convergence, the element $\theta_{z_k} \in \boldsymbol{\theta}$, $k=1, \dots, K$, is estimated by:

$$\hat{\theta}_{z_k} = \frac{n_{z_k} + \alpha_{z_k}}{\sum_z (n_z + \alpha_z)}, \quad (1)$$

where n_z is the number of times that a visual word is assigned to activity z_k in the observation.

The incorporation of topic models into cognitive modeling has several important advantages. First, as a probabilistic reasoning approach, it serves as a bridge to allow information to flow from the perception module to the decision making module. Second, the ability to model per-observation activity distribution allows topic models to take into account the risks of all robot actions in a probabilistic way and make an appropriate decision. Third, by introducing an extrinsic evaluation metric for topic model selection, the constructed cognitive system is able to accurately interpret human activities. Fourth, the unsupervised nature of topic modeling, which is explored using our new intrinsic metric, facilitates online discovery of new knowledge (e.g., human activities). All these advantages allow us to apply topic models to construct an artificial cognitive system that is able to better interpret human activities, discover new knowledge and react more appropriately and safely to humans, which is highly desirable for real-world online human-robot interaction scenarios.

4. Interpretability and Generalizability

To improve artificial cognitive modeling, we introduce two novel indicators and discuss their relationship in this section, which are the core of the *Self-Reflection* module in Fig. 1.

4.1. Interpretability Indicator

We observe that accuracy is not an appropriate assessment metric for robot decision making, since it only considers the most probable human activity category and ignores the others. To utilize the category distribution, which contains much richer information, the *interpretability indicator*, denoted by I_I , is introduced. I_I is able to encode how well topic modeling matches human common sense. Like the accuracy metric, I_I is an extrinsic metric, meaning that it requires a ground truth to compute. Formally, I_I is defined as follows:

Definition 1 (Interpretability indicator). Given the observation w with the ground truth g and the distribution θ over $K \geq 2$ categories, let $\theta_s = (\theta_1, \dots, \theta_{k-1}, \theta_k, \theta_{k+1}, \dots, \theta_K)$ denote the sorted proportion satisfying $\theta_1 \geq \dots \geq \theta_{k-1} \geq \theta_k \geq \theta_{k+1} \geq \dots \geq \theta_K \geq 0$ and $\sum_{i=1}^K \theta_i = 1$, and let $k \in \{1, \dots, K\}$ represent the index of the assignment in θ_s that matches g . The interpretability indicator $I_I(\theta, g) = I_I(\theta_s, k)$ is defined as:

$$I_I(\theta_s, k) \triangleq \frac{1}{a} \left(\frac{K-k}{K-1} + \mathbb{1}(k=K) \right) \left(\frac{\theta_k}{\theta_1} - \frac{\theta_{k+\mathbb{1}(k \neq K)}}{\theta_k} + b \right) \quad (2)$$

where $\mathbb{1}(\cdot)$ is the indicator function, and $a = 2$, $b = 1$ are normalizing constants.

The indicator I_I is defined over the per-observation category proportion θ , which takes values in the $(K-1)$ -simplex [51]. The sorted proportion θ_s is computed through sorting θ , which is inferred by topic models. In the definition, the ground truth is represented by its location in θ_s , i.e., the k -th most probable assignment in θ_s matches the ground truth label. The indicator function $\mathbb{1}(\cdot)$ in Eq. (2) is adopted to deal with the special case when $k = K$.

For an observation in an activity recognition task with K categories, given its ground truth index k and sorted category proportion θ_s , we summarize I_I 's properties as follows:

Proposition 1 (I_I 's properties). The interpretability indicator $I_I(\theta, g) = I_I(\theta_s, k)$ satisfies the following properties:

1. If $k = 1$, $\forall \theta_s$, $I_I(\theta_s, k) \geq 0.5$.
2. If $k = K$, $\forall \theta_s$, $I_I(\theta_s, k) \leq 0.5$.
3. $\forall \theta_s$, $I_I(\theta_s, k) \in [0, 1]$.
4. $\forall k \in \{1, \dots, K\}$ and θ_s, θ'_s such that $\theta_1 \geq \theta'_1$, $\theta_k = \theta'_k$ and $\theta_{k+\mathbb{1}(k \neq K)} = \theta'_{k+\mathbb{1}(k \neq K)}$, $I_I(\theta_s, k) \leq I_I(\theta'_s, k)$ holds.
5. $\forall k \in \{1, \dots, K\}$ and θ_s, θ'_s such that $\theta_{k+\mathbb{1}(k \neq K)} \geq \theta'_{k+\mathbb{1}(k \neq K)}$, $\theta_1 = \theta'_1$ and $\theta_k = \theta'_k$, $I_I(\theta_s, k) \leq I_I(\theta'_s, k)$ holds.
6. $\forall k \in \{1, \dots, K\}$ and θ_s, θ'_s such that $\theta_k \geq \theta'_k$, $\theta_1 = \theta'_1$ and $\theta_{k+\mathbb{1}(k \neq K)} = \theta'_{k+\mathbb{1}(k \neq K)}$, $I_I(\theta_s, k) \geq I_I(\theta'_s, k)$ holds.

7. $\forall k, k' \in \{1, \dots, K\}$ such that $k \leq k' < K$ and $\forall \theta_s, \theta'_s$ such that $\theta_k = \theta'_k$, $\theta_1 = \theta'_1$ and $\theta_{k+1}(k \neq K) = \theta'_{k+1}(k \neq K)$, $I_I(\theta_s, k) \geq I_I(\theta'_s, k')$ holds.

PROOF. See Appendix A.

The indicator I_I is able to quantitatively measure how well topic modeling can match human common sense, because it captures three essential considerations to simulate the process of how humans evaluate the category proportion θ :

- A topic model performs better, in general, if it obtains a larger θ_k (Property 6). In addition, a larger θ_k generally indicates θ_k is closer to the beginning in θ_s and further away from the end (Property 7).

Example: A topic model obtaining the sorted proportion $[0.4, \boxed{0.35}, 0.15, 0.10]$ performs better than a model obtaining $[0.4, \boxed{0.30}, 0.15, 0.15]$, where the ground truth is marked with a box, i.e., $k = 2$ in the example.

- A smaller difference between θ_k and θ_1 indicates better modeling performance (Properties 4 and 5), in general. Since the resulting category proportion is sorted, a small difference between θ_k and θ_1 guarantees θ_k has an even smaller difference from θ_2 to θ_{k-1} .

Example: A topic model obtaining the sorted proportion $[0.4, \boxed{0.3}, 0.2, 0.1]$ performs better than the model with the proportion $[0.5, \boxed{0.3}, 0.2, 0]$.

- A larger distinction between θ_k and θ_{k+1} generally indicates better modeling performance (Properties 5 and 6), since it better separates the correct assignment from the incorrect assignments with lower probabilities.

Example: A topic model obtaining the sorted proportion $[0.4, \boxed{0.4}, 0.1, 0.1]$ performs better than the topic model obtaining the proportion $[0.4, \boxed{0.4}, 0.2, 0]$.

The indicator I_I extends the *accuracy* metric I_A (i.e., rate of correctly recognized data), as described in Proposition 2:

Proposition 2 (Relationship of I_I and I_A). *The accuracy measure I_A is a special case of $I_I(\theta_s, k)$, when $\theta_1 = 1.0$, $\theta_2 = \dots = \theta_K = 0$, and $k = 1$ or $k = K$.*

PROOF. See Appendix B.

4.2. Generalizability Indicator

An artificial cognitive model requires the crucial capability of detecting new situations and being aware that the learned knowledge becomes less applicable in an online fashion. To this end, we propose the *generalizability indicator* (I_G), an intrinsic metric that does not require ground truth to compute and consequently can be used online.

The introduction of I_G is inspired by the perplexity metric (also referred to as held-out likelihood), which evaluates a topic model’s generalization ability on a fraction of held-out instances using cross-validation [54] or unseen observations [55]. The perplexity is defined as the log-likelihood of words in an observation [32]. Because different observations may contain a different number of visual words, we compute the *Per-Visual-Word Perplexity* ($Pvwp$). Mathematically, given the trained topic model \mathcal{M} and an observation \mathbf{w} , $Pvwp$ is defined as follows:

$$Pvwp(\mathbf{w}|\mathcal{M}) = \frac{1}{N} \log P(\mathbf{w}|\mathcal{M}) = \frac{1}{N} \log \prod_{n=1}^N P(w_n|\mathbf{w}_{<n}, \mathcal{M}) \quad (3)$$

where $N = |\mathbf{w}|$ is the number of visual words in \mathbf{w} , and the subscript $< n$ denotes positions before n . Because $P(\mathbf{w}|\mathcal{M})$ is a probability that satisfies $P(\mathbf{w}|\mathcal{M}) \leq 1$, it is guaranteed $Pvwp(\mathbf{w}|\mathcal{M}) \leq 0$. The left-to-right algorithm, presented in Algorithm 1, is used to estimate $Pvwp$, which is an accurate and efficient Gibbs sampling method to estimate perplexity [32]. The algorithm decomposes $P(\mathbf{w}|\mathcal{M})$ in an incremental, left-to-right fashion, where the subscript $\neg n$ is a quantity that excludes data from the n th position. Given observations $\mathcal{W} = \{\mathbf{w}_1, \dots, \mathbf{w}_M\}$, $Pvwp(\mathcal{W}|\mathcal{M})$ is defined as the average of each observation’s perplexity:

$$Pvwp(\mathcal{W}|\mathcal{M}) = \frac{1}{M} \sum_{m=1}^M Pvwp(\mathbf{w}_m|\mathcal{M}) \quad (4)$$

Based on $Pvwp$, the generalizability indicator I_G , on previously unseen observations in the testing phase or using the held-out instances in cross-validation, is defined as follows:

Definition 2 (Generalizability indicator). *Let \mathcal{M} denote a trained topic model, \mathcal{W}_{valid} denote the validation dataset that is used in the training phase, and \mathbf{w} be an previously*

Algorithm 1: Left-to-right $Pvwp$ estimation

Input : \mathbf{w} (observation), \mathcal{M} (trained topic model), and R (number of particles)

Output : $Pvwp(\mathbf{w}|\mathcal{M})$

```
1: Initialize  $l = 0$  and  $N = |\mathbf{w}|$ ;  
2: for each position  $n = 1$  to  $N$  in  $\mathbf{w}$  do  
3:   Initialize  $p_n = 0$ ;  
4:   for each particle  $r = 1$  to  $R$  do  
5:     for  $n' < n$  do  
6:       Sample  $z_{n'}^{(r)} \sim P(z_{n'}^{(r)} | w_{n'}, \{z_{<n}^{(r)}\}_{-n'}, \mathcal{M})$ ;  
7:     end  
8:     Compute  $p_n = p_n + \sum_t P(w_n, z_n^{(r)} = t | z_{<n}^{(r)}, \mathcal{M})$ ;  
9:     Sample  $z_n^{(r)} \sim P(z_n^{(r)} | w_n, z_{<n}^{(r)}, \mathcal{M})$ ;  
10:   end  
11:   Update  $p_n = \frac{p_n}{R}$  and  $l = l + \log p_n$ ;  
12: end  
13: return  $Pvwp(\mathbf{w}|\mathcal{M}) \simeq \frac{l}{N}$ .
```

unseen observation. We define the generalizability indicator:

$$I_G(\mathbf{w}) \triangleq \begin{cases} \frac{\exp(Pvwp(\mathbf{w}|\mathcal{M}))}{c \cdot \exp(Pvwp(\mathcal{W}_{valid}|\mathcal{M}))} \\ \text{if } \exp(Pvwp(\mathbf{w}|\mathcal{M})) < c \cdot \exp(Pvwp(\mathcal{W}_{valid}|\mathcal{M})) \\ 1 \text{ if } \exp(Pvwp(\mathbf{w}|\mathcal{M})) \geq c \cdot \exp(Pvwp(\mathcal{W}_{valid}|\mathcal{M})) \end{cases} \quad (5)$$

where $c \in [1, \infty)$ is a constant encoding novelty levels.

Besides considering the topic model's generalization ability, I_G also evaluates whether previously unseen observations are well-represented by the training set, i.e., whether the training set used to train the topic model is exhaustive. The training set is defined as *exhaustive* when it contains instances from all categories that can possibly be observed in the testing phase [56]. When some categories are missing and not represented by the training set, it is defined as *non-exhaustive*; in this case, novel categories emerge in the testing phase. Since it is impractical, often impossible, to define an exhaustive training set, mainly because some of the categories may not exist at the time of training, the ability to discover novelty and be aware that the learned model is less applicable

is essential for safe, adaptive decision making. The indicator I_G provides this ability through evaluating how well new observations are represented by the validation set in the training phase.

We constrain I_G 's value in the range $(0, 1]$, with a greater value indicating less novelty, which means an observation can be better encoded by the training set and the topic model generalizes better on this observation. The constant c in Eq. (5) provides the flexibility to encode the degree to which we consider an observation to be novel.

4.3. Indicator Relationship

While the interpretability indicator interprets human activity distributions in a way that is similar to human reasoning, the generalizability indicator endows a co-robot with the self-reflection capability. We summarize their relationship in the cases when a training set is exhaustive (i.e., training contains all possible categories) and non-exhaustive (i.e., new human behavior occurs during testing), as follows:

Observation (Relationship of I_G and I_I): Let \mathcal{W}_{train} be the training dataset used to train a topic model, and I_I and I_G be the model's interpretability and generalizability indicators.

- If \mathcal{W}_{train} is exhaustive, then $I_G \rightarrow 1$ and I_I is generally independent of I_G .
- If \mathcal{W}_{train} is non-exhaustive, then I_G takes values that are much smaller than 1; I_I also takes small values and is moderately to strongly correlated with I_G .

Table 1: Meaning and relationship of I_I and I_G . The gray area denotes that the situation is generally impossible.

	I_G : low	I_G : high
I_I : low	Category is novel Model is <i>not</i> applicable	Category is <i>not</i> novel Model is <i>not</i> well interpreted
I_I : high		Category is <i>not</i> novel Model is well interpreted

This observation answers the critical question of whether a better generalized topic model can lead to better recognition performance. Intuitively, if \mathcal{W}_{train} is non-exhaustive

Table 2: Risk levels as prior knowledge to our cognitive model.

Levels	Values	Definition
Low risk	[1,30]	Unsatisfied with the robot’s performance.
Medium risk	[31,60]	Annoyed or upset by the robot’s actions.
High risk	[61,90]	Interfered with, interrupted, or obstructed.
Critical risk	[95,100]	Injured or worse (i.e., a safety risk).

and a previously unseen observation w belongs to a novel category, which is indicated by a small I_G value, a topic model trained on \mathcal{W}_{train} cannot accurately classify w . On the other hand, if w belongs to a category that is known in \mathcal{W}_{train} , then $I_G \rightarrow 1$ and the recognition performance over w only depends on the model’s performance on the validation set used in the training phase. The meaning and relationship of the indicators I_I and I_G are summarized in Table 1, where the gray area denotes that it is generally impossible for a topic model to obtain a low generalizability but a high interpretability, as a model is never correct when presented with a novel activity.

5. Self-Reflective Risk-Aware Decision Making

Another contribution of this research is a decision making framework that is capable of incorporating activity category distribution, robot self-reflection (enabled by the indicators), and co-robot action risk, which is realized in the module of *Decision Making* in Fig. 1. Our new self-reflective risk-aware decision making algorithm is presented in Algorithm 2.

Given the robot action set $\mathbf{a} = \{a_1, \dots, a_S\}$ and the human activity set $\mathbf{z} = \{z_1, \dots, z_K\}$, an action-activity risk r_{ij} is defined as the amount of discomfort, interference, or harm that can be expected to occur during the time period if the robot takes a specific action $a_i, \forall i \in \{1, \dots, S\}$ in response to an observed human activity $z_j, \forall j \in \{1, \dots, K\}$. While θ and I_G are computed online, the risks $\mathbf{r} = \{r_{ij}\}_{S \times K}$, with each element $r_{ij} \in [0, 100]$, are manually estimated off-line by domain experts and used as a prior in the decision making module. In practice, the amount of risk is

Algorithm 2: Self-reflective risk-aware decision making

Input : \mathbf{w} (observation), \mathcal{M} (trained topic model), and \mathcal{N} (decision making bipartite network)

Output: a^* (Selected robot action with minimum risk)

- 1: Estimate per-observation activity proportion θ of \mathbf{w} ;
 - 2: Compute generalizability indicator $I_G(\mathbf{w})$;
 - 3: **for** each robot action $i = 1$ to S **do**
 - 4: Estimate activity-independent risk: $r_i^{in} = \frac{1}{K} \sum_{j=1}^K r_{ij}$;
 - 5: Calculate activity-dependent risk: $r_i^{de} = \sum_{j=1}^K (\theta_j \cdot r_{ij})$;
 - 6: Combine activity-independent and dependent risks, and assign to
per-observation action risk vector: $\mathbf{r}^a(i) = (1 - I_G(\mathbf{w})) \cdot r_i^{in} + I_G(\mathbf{w}) \cdot r_i^{de}$;
 - 7: **end**
 - 8: Select optimal robot action a^* with minimum risk in \mathbf{r}^a ;
 - 9: **return** a^* .
-

categorized into a small number of risk levels for simplicity's sake. To assign a value to r_{ij} , a risk level is first selected. Then, a risk value is determined within that risk level. As listed in Table 2, we define four risk levels with different risk value ranges in our application. We intentionally leave a five-point gap between critical risk and high risk to increase the separation of critical risk from high risk actions.

A bipartite network $\mathcal{N} = \{\mathbf{a}, \mathbf{z}, \mathbf{r}\}$ is proposed to graphically illustrate the risk matrix \mathbf{r} of robot actions \mathbf{a} associated with human activities \mathbf{z} . In this network, vertices are divided into two disjoint sets \mathbf{a} and \mathbf{z} , such that every edge with a weight r_{ij} connects a vertex $a_i \in \mathbf{a}$ to a vertex $z_j \in \mathbf{z}$. An example of such a bipartite network is depicted in Fig. 2 for assistive robotics applications. The bipartite network also has a tabular representation (for example, in Table 4). Given the bipartite network, for a new observation \mathbf{w} , after θ and $I_G(\mathbf{w})$ are computed in the probabilistic reasoning module, the robot action $a^* \in \mathbf{a}$ is selected according to:

$$a^* = \arg \min_{a_i: i=1, \dots, S} \left(\frac{1 - I_G(\mathbf{w})}{K} \cdot \sum_{j=1}^K r_{ij} + I_G(\mathbf{w}) \cdot \sum_{j=1}^K (\theta_j \cdot r_{ij}) \right) \quad (6)$$




Figure 2: An illustrative example of a bipartite network (left) and the per-observation activity distribution (right) in assistive robotics applications.

The risk of taking a specific robot action is determined by two separate components: activity-independent and activity-dependent action risks. The activity-independent risk (that is $\frac{1}{K} \sum_{j=1}^K r_{ij}$) measures the inherent risk of an action, which is independent of the human activity context information, i.e., computing this risk does not require the category distribution. For example, the robot action “standing-by” generally has a smaller risk than “moving backward” in most situations. The activity-dependent risk (that is $\sum_{j=1}^K (\theta_j \cdot r_{ij})$) is the average risk weighted by context-specific information (i.e., the human activity distribution). The combination of these two risks is controlled by I_G , which automatically encodes preference over robot actions. When the learned model generalizes well over w , i.e., $I_G(w) \rightarrow 1$, the decision making process prefers co-robot actions that are more appropriate to the recognized human activity. Otherwise, if the model generalizes poorly, indicating new human activities occur and the learned model is less applicable, our decision making module would ignore the recognition results and select co-robot actions with lower activity-independent risk.

6. Experiments

6.1. Datasets and Visual Features

We employ three real-world benchmark datasets to evaluate our cognitive model on HAR tasks, which are widely used in the machine vision community: the Weizmann

activity dataset [57], the KTH activity dataset [58], and the UTK 3D activity dataset [12]. Illustrative examples from each activity category in the datasets are depicted in Fig. 3.




Figure 3: Exemplary frames of actions in the datasets used in our experiments.

In our experiments, we apply different types of local visual features to encode these datasets. For 2D datasets that contain only color videos (i.e., the Weizmann and KTH datasets), we use two different features: scale-invariant feature transform (SIFT) features [22] and space-time interest points (STIP) features [58]. For 3D datasets that contain both color and depth videos (i.e., the UTK dataset), we use the 4-dimensional local spatio-temporal features (4D-LSTF) [12].

SIFT features are the most commonly applied local visual features and have de-

sirable characteristics including invariance to transformation, rotation and scale, and robustness to partial occlusion [22]. We employ the algorithm and implementation in [22] to detect and describe SIFT features. A disadvantage of SIFT features in HAR tasks is that these features are extracted in a frame-by-frame fashion, i.e., SIFT features do not capture any temporal information. To encode time information, we also apply STIP along with the histogram of oriented gradients (HOG) and histogram of optical flow (HOF) descriptors [58]. These two types of features are extracted using only color or intensity information. Previous work has demonstrated that local features incorporating both depth and color information can greatly improve recognition accuracy [12]. Therefore, for the 3D UTK dataset we use 4D-LSTF [12] features, which are highly robust and distinct and are generated using both color and depth videos. It is also noteworthy that SIFT and STIP features can be directly extracted from color or depth videos in the 3D dataset.

These feature extraction algorithms generate a collection of feature vectors for each visual observation. Then, the feature vectors are clustered into discrete visual words using the k -means algorithm, and the number of clusters is set equal to the dictionary size. Lastly, each feature vector is indexed by a discrete word that represents cluster assignment. At this point, each observation is encoded by a BoW representation, which can be perceived by topic modeling. Although we only test the most widely used features, one should note that our artificial cognitive model is capable of incorporating different types of local visual features, since our reasoning and decision making process is independent of the features given their BoW representation.

6.2. Activity Recognition

We first evaluate the SRAC model’s capability to recognize human activities using the interpretability indicator I_I , when the training set is exhaustive. In this experiment, each dataset is split into disjoint training and testing sets. We randomly select 75% of data instances in each category as the training set, and employ the rest of the instances for testing. During training, fourfold cross-validation is used to estimate model parameters. Then, the interpretability of the topic model is computed using the testing set, which is fully represented by the training set and does not contain novel human

activities. This training-testing process is repeated five times to obtain reliable results.




Figure 4: Variations of model interpretability and its standard deviation versus dictionary size using different visual features over benchmark datasets.

Experimental results of the interpretability and its standard deviation versus the dictionary size are illustrated in Fig. 4. Our SRAC model obtains promising recognition performance in terms of interpretability: 0.989 is obtained using the STIP feature and a dictionary size 1800 on the Weizmann dataset, 0.952 using the STIP feature and a dictionary size 2000 on the KTH dataset, and 0.936 using the 4D-LSTF feature and a dictionary size 1600 on the UTK3D dataset. In general, STIP features perform better than SIFT features for color data, and 4D-LSTF features perform the best for RGB-D visual data. The dictionary size in the range [1500, 2000] can generally result in satis-




Figure 5: Model interpretability over the activities in the UTK3D dataset using different features and a dictionary size of 1600.

factory human activity recognition performance. The results are also very consistent, as illustrated by the small error bars in Fig. 4, which demonstrates our interpretability indicator’s consistency.

The model’s interpretability is also evaluated over different activity categories using the UTK3D dataset, which includes more complex activities (i.e., sequential activities) and contains more information (i.e., depth). It is observed that topic modeling’s interpretability varies for different activities. This performance variation is affected by three main factors: the topic model’s modeling ability, feature and BoW’s representability, and human activity complexity and similarity. For example, since the LDA topic model and SIFT features are not capable of modeling time, the reversal human activities including “lifting a box” and “removing a box” in the UTK3D dataset cannot be well distinguished, as illustrated in Fig. 5. Since sequential activities (e.g., “removing a box”) are more complex than repetitive activities (e.g., “waving”), they generally result in low interpretability. Since “pushing” and “walking” are similar, which share motions such as moving forward, they can also reduce interpretability. This observation provides general guidance for designing future recognition systems with the SRAC model.

6.3. Knowledge Discovery

We evaluate the SRAC model’s capability to discover new situations using the generalizability indicator I_G , when the training dataset is non-exhaustive (i.e., new human activities occur during testing). A non-exhausted setup is created by dividing the used benchmark datasets as follows. We place all data instances of one activity in the *unknown testing set*, and randomly select 25% of the instances from the remaining activities in the *known testing set*. The remaining instances are placed in the training set for learning, based on fourfold cross-validation. To evaluate the model’s ability to discover each individual human activity, given a dataset that contains K activity categories, the experiments are repeated K times, each using one category as the novel activity. Visual features that achieve the best model interpretability over each dataset are used in this set of experiments i.e., STIP features for the Weizmann and KTH datasets and 4D-LSTF features for the UTK3D dataset.

Variations of $Pvwp$ values versus the dictionary size over the validation set (in cross-validation), known testing set, and unknown testing set are shown in Fig. 6. Several important phenomena are observed. First, there exists a large $Pvwp$ gap between the known and unknown testing sets, as shown by the gray area in the figure, indicating that topic models generalize differently over data instances from known and unknown activities. A better generalization result indicates a less novel instance, which can be better represented by the training set. Since data instances from the known testing and validation sets are well represented by the training set, the $Pvwp$ gap between them is small. As shown in Fig. 6(a), it is possible that the known testing set’s $Pvwp$ value is greater than the $Pvwp$ value of the validation set, if its data instances can be better represented by the training set. Second, Fig. 6 shows that the gap’s width varies over different datasets: the Weizmann dataset generally has the largest $Pvwp$ gap, followed by the KTH dataset, and then the UTK3D dataset. The gap’s width mainly depends on the observation’s novelty, in terms of the novel activity’s similarity to the activities in the training dataset. This similarity is encoded by the portion of overlapping features. A more novel activity is generally represented by a set of more distinct visual features with less overlapping with the features existing during training, which generally results in a larger gap. For example, activities in the Weizmann dataset share fewer motions

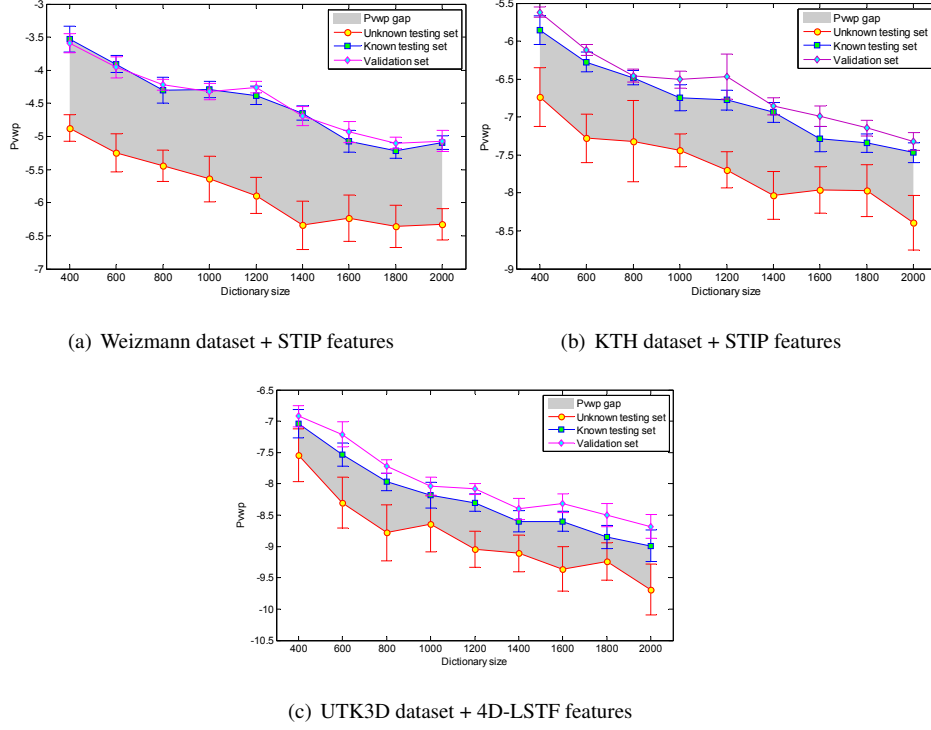


Figure 6: Variations of topic modeling’s $Pwvp$ versus dictionary size over validation set, known and unknown testing sets.

and thus contain a less number of overlapping features, which leads to a larger gap. Third, when the dictionary size increases, the model’s $Pwvp$ values decrease at a similar rate. This is because in this case, the probability of a specific codeword appearing in an instance decreases, resulting in a decreasing $Pwvp$ value.

The generalizability indicator I_G ’s characteristics are also empirically validated on the known and unknown testing sets, as illustrated in Fig. 7. An important characteristic of I_G is its invariance to dictionary size. Because $Pwvp$ over testing and validation sets has similar decreasing rate, the division operation in Eq. (5) removes the variance to dictionary size. In addition, a more novel activity generally leads to a smaller I_G value. For example, the Weizmann dataset has a smaller I_G value over the unknown testing set, because its activities are more novel in the sense that they share less overlapping motions. In general, we observe I_G is smaller than 0.5 for unknown activities




Figure 7: Variations of our model’s generalizability versus dictionary size over known and unknown testing sets for all datasets.

and greater than 0.7 for activities that are included in training sets. As indicated by the gray area in Fig. 7, similar to *Pvvp*, there exists a large gap between the I_G values over the unknown and known testing datasets. The average I_G gap across different dictionary sizes is 0.69 for the Weizmann dataset, 0.48 for the KTH dataset, and 0.36 for the UTK3D dataset. This reasoning process, based on I_G , provides a co-robot with the critical self-reflection capability, and allows a robot to reason about when new situations occur as well as when the learned model becomes less applicable.

We have pointed out that the indicator I_G is heavily affected by the novelty of an activity in terms of its proportion of overlapping features. To validate this conclusion, we generate a synthetic dataset by manually controlling the proportion of overlapping visual words in the testing instances. In order to make the characteristics of the synthetic dataset as close as possible to real-world datasets, features used in the simulation are




Figure 8: Variations of our model generalizability versus percentage of overlapping features in synthetic data.

borrowed from the KTH dataset. Instances of two activities (i.e., “bending” and “waving2”) are used to train a topic model, which is then applied as a classifier to perform recognition in this experiment. This topic model is also applied to generate overlapping visual words for a testing instance. Another topic model, whose parameters are learned using activities “siding” and “jacking”, is used to generate non-overlapping words in the testing instance. A dictionary of size 1800 is adopted, which is created using the visual words of the KTH dataset. The used four activities contain 906 unique visual words, with each pair of activities sharing less than 1% overlapping words. We generate 50 instances for each testing set, with the number of words in each instance set to 112, which is the average number of visual words in real-world instances. We present the results of five simulations in Fig. 8, which clearly shows that, in general, I_G ’s value over testing instances increases linearly with the percentage of features that overlap with the features of known activities in the training set.

6.4. Relationship of I_G and I_I

Here, we empirically analyze the relationship between the interpretability and generalizability indicators. We first validate the correlation of I_I and I_G . In addition, we investigate additional relationships of I_I and I_G , such as the probability that $I_I \leq I_G$.

While we are able to employ the exhaustive experimental setup from Section 6.2

to analyze I_G and I_I 's relationship when testing instances are fully represented by the training set, unfortunately, we cannot use the non-exhaustive setup in Section 6.3 to validate this relationship in cases where I_G takes small values. This is because ground truth cannot be assigned to instances belonging to novel activities to compute I_I , since these activities only exist in the testing set and are not presented to our model during the training phase. Inspired by the method used to generate synthetic data in Section 6.3, we adopt a semi-exhaustive experimental setup by replacing certain portions of words in each testing instance with visual words from novel activities. This experimental setup is used to validate the indicators' relationship when the training set cannot fully represent testing instances.

Each experiment is performed using F folds, where F is the number of activities in a dataset. In each fold, we take all instances of one activity out from the dataset, which is treated as a novel activity that is not presented to the topic model in the learning phase. Then, we randomly select 75% of the instances of the remaining activities as training set, which is further divided into training and validation sets to perform four-fold cross-validation. The rest of the instances are used as an "initial" testing set. During the testing phase, the novel activity's word distribution is used to generate new visual words to replace a proportion of the words in each instance in the initial testing set. This testing is performed six times within each of the F folds using different replacement rates (i.e., 0.25, 0.35, \dots , 0.75). Testing results from all F folds are used to investigate I_I and I_G 's relationship. In this experimental setup, we use features that achieve the best interpretability over each dataset. In addition, we set the dictionary size to 1600, which achieves the best interpretability over all datasets in general.

This experimental setup is semi-exhaustive in the sense that, although training data cannot fully represent testing instances due to the replaced features that are generated from unknown activities, the remaining non-replaced features are presented to the model during the learning phase, and the ground truth assigned to each testing instance remains the same, which is also known to the model. It is noteworthy that we do not use very high or very low replacement rates. A very low replacement rate makes the experimental setup equivalent to the exhaustive setup. When using a very high replacement rate, testing instances can be viewed as being drawn from the novel activity; in

Table 3: Relationship between I_I and I_G over exhaustive and semi-exhaustive datasets.

Dataset + Features	Exhaustive		Semi-exhaustive	
	$\rho_{I,G}$	$P_{I_I \leq I_G}$	$\rho_{I,G}$	$P_{I_I \leq I_G}$
Weizmann + STIP	-0.065	0.456	0.664	0.912
KTH + STIP	0.036	0.324	0.685	0.853
UTK3D + 4D-LSTF	0.097	0.275	0.714	0.896

this case the ground truth associated with a testing instance would be meaningless or incorrect.

We empirically analyze the correlation between I_I and I_G , using both exhaustive and semi-exhaustive datasets, in order to determine whether better generalizability indicates better interpretability. The Pearson correlation coefficient is used to measure the strength and direction of the linear relationship between these two indicators. Given a dataset $\mathcal{W} = \{\mathbf{w}_1, \dots, \mathbf{w}_{|\mathcal{W}|}\}$ and its ground truth $\mathbf{g} = \{g_1, \dots, g_{|\mathcal{W}|}\}$, this correlation is mathematically defined as follows:

$$\rho_{I,G} = \frac{E[(\mathbf{I}_I - \mu_{I_I})(\mathbf{I}_G - \mu_{I_G})]}{\sigma_{I_I} \sigma_{I_G}}, \quad (7)$$

where $\mathbf{I}_G = \{I_G(\mathbf{w}_1), \dots, I_G(\mathbf{w}_{|\mathcal{W}|})\}$ and $\mathbf{I}_I = \{I_I(\mathbf{w}_1, g_1), \dots, I_I(\mathbf{w}_{|\mathcal{W}|}, g_{|\mathcal{W}|})\}$ are vectors of interpretability and generalizability indicators for all of the instances in the dataset, μ is the mean and σ is the standard deviation of the indicators in the vector. Our experimental results are listed in Table 3. For the exhaustive dataset, topic models which perform better on generalizability are not necessarily better interpreted, which is indicated by the weak linear correlation between the indicators. This is because, when testing on exhaustive datasets, I_G takes values closer to 1. But the model’s interpretability takes a wide range of values, depending on the model’s modeling capacity, feature representability and dataset complexity, as explained in Section 6.2. For semi-exhaustive datasets, I_I and I_G are moderately to strongly correlated, which indicates that a poor generalizability usually leads to a poor interpretability. Since I_G reflects the novelty of an instance as discussed in Section 6.3, a low I_G ’s value means the instance is badly represented by the training set. Therefore, the trained model cannot obtain a

good interpretability over the instance of an activity that is not well represented during the training phase.

We also check an additional relationship, i.e., the probability that I_I is smaller than or equal to I_G . Given a labeled dataset $\mathcal{W} = \{\mathbf{w}_1, \dots, \mathbf{w}_M\}$ and its ground truth $\mathbf{g} = \{g_1, \dots, g_M\}$, this probability is defined as follows:

$$P_{I_I \leq I_G} = \frac{1}{M} \sum_{m=1}^M \mathbb{1}(I_I(\mathbf{w}_m, g_m) \leq I_G(\mathbf{w}_m)). \quad (8)$$

The experimental results are presented in Table 3. One of the most important observations is that, for a majority of testing instances (more than 85%) in the semi-exhaustive experiment, I_G 's value is greater than I_I 's value. This again shows that a poor generalizability usually indicates a poor interpretability. Using the exhaustive experimental setup, it is more probable that I_G takes smaller values than I_I . This is because when the training set is exhaustive, the topic model is well trained and can well recognize testing instances, which leads to $I_I \rightarrow 1$ for most of testing instances. On the other hand, although I_G also takes a large value in general, it is usually slightly smaller than one, because features in testing instances usually do not completely overlap with features in training instances.

6.5. Decision Making

We assess our SRAC model's decision making capability using a Turtlebot 2 robot in a human following task, which is important in many human-robot teaming applications. In this task, a robotic follower needs to decide at what distance to follow the human teammate. We are interested in three human behaviors: "walking" in a straight line, "turning", and "falling", shown in Fig. 9. With perfect perception and reasoning, i.e., a robot always perfectly interprets human activities, we assume the ideal robot actions are to "stay far from the human" when he or she is walking in a straight line (to not interrupt the human), "move close to the human" when the subject is turning (to avoid losing the target), and "stop beside the human" when he or she is falling (to provide assistance).

In order to qualitatively assess the performance, we collect 20 color-depth instances from each human behaviors to train the SRAC model, using a BoW representation




Figure 9: Experiment setup for validating the SRAC model’s decision making ability in a human following task using a Turtlebot 2 robot.

Table 4: The risk matrix used in the robot following task.

Robot Actions	Falling	Turning	Walking
Stay besides humans	0	20	50
Move close	90	0	20
Stay far away	95	80	0

based on 4D-LSTF features. The risk matrix used in this task is presented in Table 4. We evaluate our model in two circumstances. *Case 1*: exhaustive training (i.e., no unseen human behaviors occur in testing). In this case, the subjects only perform the three activities during testing with small variations in motion speed and style. *Case 2*: non-exhaustive training (i.e., novel movements occur during testing). In this case, the subjects not only perform the activities with large variations, but also add additional movements (such as jumping and squatting) which are not observed in the training phase. During testing, each activity is performed 40 times. The model performance is measured using failure rate, i.e., the percentage with which the robot fails to stop besides to help the human or loses the target.

Experimental results are presented in Table 5, where the traditional methodology, which selects the co-robot actions only based on the most probable human activity, is used as a baseline for comparison. We observe that the proposed SRAC model significantly decreases the failure rate in both exhaustive and non-exhaustive setups. When

the training set is exhaustive and no new activities occur during testing (Case 1), the results demonstrate that incorporating human activity distributions and robot action risks improves decision making performance. When the training set is non-exhaustive and new activities occur during testing (Case 2), the SRAC model significantly outperforms the baseline model. In this situation, if I_G has a very small value, according to Eq. 6, our model tends to select safer robot actions, i.e., “stay beside humans,” since its average risk is the lowest, which is similar to the human common practice “playing it safe in uncertain times.” The results show the importance of self-reflection for decision making especially under uncertainty.

Table 5: Failure rate (%) in exhaustive (Case 1) and non-exhaustive (Case 2) experimental settings.

Exp. settings	Models	Fail to assist	Fail to follow
Exhaustive (Case 1)	Baseline	10.5%	15%
	SRAC	0.5%	5.5%
Non-exhaustive (Case 2)	Baseline	45.5%	60%
	SRAC	24.5%	35.5%

7. Conclusion

In this paper, we construct an artificial cognitive model that provides co-robots with both accurate perception and new information discovery capabilities, which enables safe, reliable robot decision making for HAR tasks in human-robot interaction applications.

The proposed SRAC model exploits topic modeling, which is unsupervised and allows for the discovery of new knowledge without presence in training. In addition, topic modeling is also able to treat activity estimation as a distribution and incorporate risks for each action response, which is beneficial for the system’s ability to make decisions. In order to provide the capability of accurate human activity interpretation, we define a new interpretability indicator (I_I) and demonstrate its ability to enable a robot to interpret category distribution in a similar fashion to humans. The indicator I_I is applied to map detected clusters to known activity categories and select the best

interpreted model. In addition, to provide the ability of knowledge discovery, we introduce a novel generalizability indicator (I_G). It measures how well an observation can be represented by the learned knowledge, which allows for self-reflection that can enable the SRAC model to identify new scenarios.

We applied the proposed SRAC model in extensive experiments to demonstrate the effectiveness using both synthetic and real-world datasets. We show that our model performs extremely well in terms of interpretability; that is, our model’s recognition results closely and consistently match human common sense. We demonstrate that, using I_G , our cognitive model is capable of discovering new knowledge, i.e., observations from new activity categories that are not considered in the training phase can be automatically detected. We also examine the relationship between I_I and I_G and show, both analytically and experimentally, that I_G can also be used as an indicator for I_I . The results reveal that scenarios with a low I_G score for an observation will equate to a low I_I score with high confidence, i.e., a badly generalized model is likely to be inaccurate. We further demonstrate the advantages of using distributions over activity categories, as well as the importance of the evaluation metrics in order to create a system capable of safe, reliable decision making.

Appendix A. Proof of I_I ’s Properties (Proposition 1)

If denominator in Definition 1 is 0, then limit is used. Given the normalizing constants $a = 2$ and $b = 1$:

1. If $k = 1$, $I_I(\boldsymbol{\theta}_s, k) = \frac{1}{a} \left(1 + b - \frac{\theta_2}{\theta_1} \right) = 1 - \frac{\theta_2}{2\theta_1}$. Since $\boldsymbol{\theta}_s$ is decreasingly sorted, satisfying $\theta_1 \geq \theta_2 \geq 0$, then $-\frac{\theta_2}{2\theta_1} \geq -0.5$. Thus, $I_I(\boldsymbol{\theta}_s, k) = 1 - \frac{\theta_2}{2\theta_1} \geq 0.5$

2. If $k = K$, $I_I(\boldsymbol{\theta}_s, k) = \frac{1}{a} \left(\frac{\theta_K}{\theta_1} - 1 + b \right) = \frac{\theta_K}{2\theta_1}$. Since $\boldsymbol{\theta}_s$ is decreasingly sorted, satisfying $\theta_1 \geq \theta_K \geq 0$, then $\frac{\theta_K}{\theta_1} \leq 1$. Thus, $I_I(\boldsymbol{\theta}_s, k) = \frac{\theta_K}{2\theta_1} \leq 0.5$.

3. First, we prove $I_I(\boldsymbol{\theta}_s, k) \geq 0$. Since $K \geq k > 0$ and $K \geq 2$, the second multiplier $F_2 = \frac{K-k}{K-1} + 1(k = K) > 0$. Given $b = 1$, the third multiplier satisfies $F_3 = \frac{\theta_k}{\theta_1} - \frac{\theta_{k+1}(k \neq K)}{\theta_k} + b = \frac{\theta_k^2 + \theta_1(\theta_k - \theta_{k+1}(k \neq K))}{\theta_1 \theta_k}$. Since $\boldsymbol{\theta}_s$ is decreasingly sorted, then $\theta_1 \geq \theta_k \geq \theta_{k+1}(k=K) \geq 0$. Thus, $F_3 \geq 0$. Equality is obtained when $\theta_k = \theta_{k+1}(k=K) = 0$. Since $a > 0$, $F_2 > 0$ and $F_3 \geq 0$, then $I_I(\boldsymbol{\theta}_s, k) = \frac{1}{a} \cdot F_2 \cdot F_3 \geq 0$.

Now, we prove $I_I(\boldsymbol{\theta}_s, k) \leq 1$. When $k = K$, by property 2, $I_I(\boldsymbol{\theta}_s, k) \leq 1$ directly holds. If $K > k \geq 1$, then $F_2 = \frac{K-k}{K-1} \leq 1$. Equality holds when $k = 1$. Since $\boldsymbol{\theta}_s$ is decreasingly sorted, satisfying $\theta_1 \geq \theta_k \geq \theta_{k+1} \geq 0$, then $\frac{\theta_k}{\theta_1} \leq 1$ and $\frac{\theta_{k+1}}{\theta_k} \geq 0$. Given $b = 1$, we have $F_3 = \frac{\theta_k}{\theta_1} - \frac{\theta_{k+1}}{\theta_k} + b \leq \frac{\theta_k}{\theta_1} + 1 \leq 2$. Equality holds when $\theta_k = \theta_1$ and $\theta_{k+1} = 0$. Thus, given $a = 2$, we obtain $I_I(\boldsymbol{\theta}_s, k) = \frac{1}{a} \cdot F_2 \cdot F_3 \leq 1$. Thus, $\forall \boldsymbol{\theta}$, $I_I(\boldsymbol{\theta}_s, k) \in [0, 1]$ holds.

4. Since $\forall k \in \{1, \dots, K\}$, $\boldsymbol{\theta}_s, \boldsymbol{\theta}'_s$ satisfy $\theta_k = \theta'_k$ and $\theta_{k+l(k=K)} = \theta'_{k+l(k=K)}$, we obtain $I_I(\boldsymbol{\theta}_s, k) - I_I(\boldsymbol{\theta}'_s, k) = \frac{1}{a} \left(\frac{K-k}{K-1} + l(k=K) \right) \left(\frac{\theta_k}{\theta_1 \theta'_1} (\theta'_1 - \theta_1) \right)$. Since $\frac{K-k}{K-1} + l(k=K) > 0$ and $\theta'_1 \geq \theta_1$, Then, $I_I(\boldsymbol{\theta}_s, k) - I_I(\boldsymbol{\theta}'_s, k) \leq 0$. Equality holds if $\theta'_1 = \theta_1$ or $\theta_k = 0$. Thus, $I_I(\boldsymbol{\theta}_s, k) \leq I_I(\boldsymbol{\theta}'_s, k)$.

5. Since $\forall k \in \{1, \dots, K\}$, $\boldsymbol{\theta}_s, \boldsymbol{\theta}'_s$ satisfy $\theta_1 = \theta'_1$ and $\theta_k = \theta'_k$, we obtain $I_I(\boldsymbol{\theta}_s, k) - I_I(\boldsymbol{\theta}'_s, k) = \frac{1}{a} \left(\frac{K-k}{K-1} + l(k=K) \right) \left(\frac{1}{\theta_k} (\theta'_{k+l(k=K)} - \theta_{k+l(k=K)}) \right)$. Since $\frac{K-k}{K-1} + l(k=K) > 0$ and $\theta_{k+l(k=K)} \geq \theta'_{k+l(k=K)}$, Then, $I_I(\boldsymbol{\theta}_s, k) - I_I(\boldsymbol{\theta}'_s, k) \leq 0$. Equality holds if $\theta_{k+l(k=K)} = \theta'_{k+l(k=K)}$. Thus, $I_I(\boldsymbol{\theta}_s, k) \leq I_I(\boldsymbol{\theta}'_s, k)$ holds.

6. Since $\forall k \in \{1, \dots, K\}$, $\boldsymbol{\theta}_s, \boldsymbol{\theta}'_s$ satisfy $\theta_k = \theta'_k$ and $\theta_{k+l(k=K)} = \theta'_{k+l(k=K)}$, we obtain $I_I(\boldsymbol{\theta}_s, k) - I_I(\boldsymbol{\theta}'_s, k) = \frac{1}{a} \left(\frac{K-k}{K-1} + l(k=K) \right) \left(\frac{1}{\theta_1} + \frac{\theta_{k+l(k=K)}}{\theta_k \theta'_k} \right) (\theta_k - \theta'_k)$. Since $\frac{K-k}{K-1} + l(k=K) > 0$ and $\frac{1}{\theta_1} + \frac{\theta_{k+l(k=K)}}{\theta_k \theta'_k} \geq 0$, and $\theta_k > \theta'_k$, Then, $I_I(\boldsymbol{\theta}_s, k) - I_I(\boldsymbol{\theta}'_s, k) \geq 0$. Equality holds if $\theta_k = \theta'_k$. Thus, $I_I(\boldsymbol{\theta}_s, k) \geq I_I(\boldsymbol{\theta}'_s, k)$ holds.

7. $\forall k, k' \in \{1, \dots, K\}$ satisfying $k \leq k' < K$, and $\forall \boldsymbol{\theta}_s, \boldsymbol{\theta}'_s$ satisfying $\theta_{k+1} = \theta'_{k'+1}, \theta_1 = \theta'_1$ and $\theta_k = \theta'_k$, we obtain $I_I(\boldsymbol{\theta}_s, k) - I_I(\boldsymbol{\theta}'_s, k') = \frac{1}{a(K-1)} \left(\frac{\theta_k}{\theta_1} - \frac{\theta_{k+1}}{\theta_k} + b \right) (k' - k)$. Since $K > 1$, $\frac{\theta_k}{\theta_1} - \frac{\theta_{k+1}}{\theta_k} + b \geq 0$, and $k' \geq k$, $I_I(\boldsymbol{\theta}_s, k) - I_I(\boldsymbol{\theta}'_s, k') \geq 0$, with equality holding when $\theta_k = \theta_{k+1} = 0$ or $k = k'$. Thus, $I_I(\boldsymbol{\theta}_s, k) \geq I_I(\boldsymbol{\theta}'_s, k')$ holds.

Appendix B. Proof of the Relationship between I_I and I_A (Proposition 2)

Given an observation \mathbf{w} , the accuracy metric I_A indicates whether the recognition result $y(\mathbf{w})$ matches the ground truth g . Formally, I_A is defined as follows:

$$I_A(y(\mathbf{w}), g) = l(y(\mathbf{w}) = g). \quad (\text{B.1})$$

With this definition, we prove that I_A is a special case of our I_I indicator in Definition 1 when $\theta_1 = 1.0, \theta_2 = \dots = \theta_K = 0$, and $k = 1$ or $k = K$.

Given the normalizing constants $a = 2$ and $b = 1$, when $\theta_1 = 1.0$, $\theta_2 = \dots = \theta_K = 0$, and $k = 1$ (i.e., the recognition result $y(\mathbf{w})$ matches the ground truth g), we obtain:

$$I_s(\boldsymbol{\theta}_s, 1) = \frac{1}{a} \left(\frac{K-1}{K-1} + 0 \right) \left(\frac{\theta_1}{\theta_1} - \frac{\theta_2}{\theta_1} + b \right) = \frac{b+1}{a} = 1.$$

When $k = K$ (i.e., $y(\mathbf{w}) \neq g$), we obtain:

$$I_s(\boldsymbol{\theta}_s, K) = \frac{1}{a} \left(\frac{K-1}{K-1} + 1 \right) \left(\frac{\theta_K}{\theta_1} - \frac{\theta_K}{\theta_K} + b \right) = \frac{2(b-1)}{a} = 0.$$

Combining both cases, we obtain:

$$\begin{aligned} I_s(\boldsymbol{\theta}_s, k) &= \begin{cases} 1 & \text{if } k = 1 \text{ (i.e., } y(\mathbf{w}) = g) \\ 0 & \text{if } k = K \text{ (i.e., } y(\mathbf{w}) \neq g) \end{cases} \\ &= l(y(\mathbf{w}) = g). \end{aligned} \quad (\text{B.2})$$

We observe Eq. (B.1) is equivalent to Eq. (B.2), and thereby prove that I_A is a special case of the I_I indicator in the cases when $\theta_1 = 1.0$, $\theta_2 = \dots = \theta_K = 0$, and $k = 1$ or $k = K$.

References

- [1] F. J. Varela, J. Dupuy, Understanding Origins, Kluwer Academic Publishers, 1992.
- [2] D. Vernon, G. Metta, G. Sandini, A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents, IEEE Transactions on Evolutionary Computation 11 (2) (2007) 151–180.
- [3] J. R. Anderson, ACT: A simple theory of complex cognition, American Psychologist 51 (1996) 355–365.
- [4] G. Trafton, L. Hiatt, A. Harrison, F. Tamborello, S. Khemlani, A. Schultz, ACT-R/E: An embodied cognitive architecture for human-robot interaction, Journal of Human-Robot Interaction 2 (1) (2013) 30–55.

- [5] C. L. Dancy, ACT-R Φ : A cognitive architecture with physiology and affect, *Biologically Inspired Cognitive Architectures* 6 (2013) 40–45.
- [6] J. Laird, *The Soar cognitive architecture*, MIT Press, 2012.
- [7] D. Isla, R. Burke, M. Downie, B. Blumberg, A layered brain architecture for synthetic creatures, in: *International Joint Conferences on Artificial Intelligence*, 2001.
- [8] C. Burghart, R. Mikut, R. Stiefelhagen, T. Asfour, H. Holzapfel, P. Steinhaus, R. Dillmann, A cognitive architecture for a humanoid robot: a first approach, in: *IEEE-RAS International Conference on Humanoid Robots*, 2005.
- [9] U. Schmid, M. Ragni, C. Gonzalez, J. Funke, The challenge of complexity for cognitive systems, *Cognitive Systems Research* 12 (3-4) (2011) 211–218.
- [10] Q. V. Le, W. Y. Zou, S. Y. Yeung, A. Y. Ng, Learning hierarchical invariant spatio-temporal features for action recognition with independent subspace analysis, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2011.
- [11] A. Alahi, R. Ortiz, P. Vandergheynst, FREAK: Fast retina keypoint, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.
- [12] H. Zhang, L. E. Parker, 4-dimensional local spatio-temporal features for human activity recognition., in: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011.
- [13] Y. Girdhar, P. Giguere, G. Dudek, Autonomous adaptive exploration using realtime online spatiotemporal topic modeling, *International Journal of Robotics Research* 33 (4) (2013) 645–657.
- [14] J. C. Niebles, H. Wang, L. Fei-Fei, Unsupervised learning of human action categories using spatial-temporal words, *International Journal of Computer Vision* 79 (3) (2008) 299–318.
- [15] J. Aggarwal, M. Ryoo, Human activity analysis: A review, *ACM Computing Surveys* 43 (3) (2011) 16:1–16:43.

- [16] P. V. K. Borges, N. Conci, A. Cavallaro, Video-based human behavior understanding: a survey, *IEEE Transactions on Circuits and Systems for Video Technology* 23 (11) (2013) 1993–2008.
- [17] X. Chen, T. Yang, Extraction method of gait feature based on human centroid trajectory, in: *Computer Engineering and Networking*, 2014, pp. 515–523.
- [18] W. Ge, R. T. Collins, R. B. Ruback, Vision-based analysis of small groups in pedestrian crowds, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34 (5) (2012) 1003–1016.
- [19] M. Singh, A. Basu, M. Mandal, Human activity recognition based on silhouette directionality, *IEEE Transactions on Circuits and Systems for Video Technology* 18 (9) (2008) 1280–1292.
- [20] I. N. Junejo, K. N. Junejo, Z. Al Aghbari, Silhouette-based human action recognition using SAX-Shapes, *The Visual Computer* 30 (3) (2014) 259–269.
- [21] H. Zhang, L. E. Parker, Bio-inspired predictive orientation decomposition of skeleton trajectories for real-time human activity prediction, in: *IEEE International Conference on Robotics and Automation*, 2015.
- [22] D. G. Lowe, Distinctive image features from scale-invariant keypoints, *International Journal of Computer Vision* 60 (2) (2004) 91–110.
- [23] X. Sun, M. Y. Chen, A. Hauptmann, Action recognition via local descriptors and holistic features, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [24] A. Behera, A. G. Cohn, D. C. Hogg, Real-time activity recognition by discerning qualitative relationships between randomly chosen visual features, in: *British Machine Vision Conference*, 2014.
- [25] C. Schuldt, I. Laptev, B. Caputo, Recognizing human actions: A local SVM approach, in: *International Conference on Pattern Recognition*, 2004.

- [26] P. Dollár, V. Rabaud, G. Cottrell, S. Belongie, Behavior recognition via sparse spatio-temporal features, in: IEEE International Workshop on Visual Surveillance and Performance Evaluation of Tracking and Surveillance, 2005.
- [27] F. Zhao, Y. Huang, L. Wang, T. Tan, Relevance topic model for unstructured social group activity recognition, in: Advances in Neural Information Processing Systems, 2013.
- [28] Y. Wang, G. Mori, Human action recognition by semilattent topic models, IEEE Transactions on Pattern Analysis and Machine Intelligence 31 (10) (2009) 1762–1774.
- [29] T. Huynh, M. Fritz, B. Schiele, Discovery of activity patterns using topic models, in: International Conference on Ubiquitous Computing, 2008.
- [30] K. Farrahi, D. Gatica-Perez, Discovering routines from large-scale human locations using probabilistic topic models, ACM Transactions on Intelligent Systems and Technology 2 (1) (2011) 3:1–3:27.
- [31] R. G. Freedman, H.-T. Jung, S. Zilberstein, Temporal and object relations in plan and activity recognition for robots using topic models, in: AAAI Fall Symposium Series, 2014.
- [32] H. Wallach, I. Murray, R. Salakhutdinov, D. Mimno, Evaluation methods for topic models, in: International Conference on Machine Learning, 2009.
- [33] X. Wei, W. B. Croft, LDA-based document models for ad-hoc retrieval, in: Iterational Conference on Research and Development in Information Retrieval, 2006.
- [34] J. Chang, J. Boyd-Graber, S. Gerrish, C. Wang, D. Blei, Reading tea leaves: How humans interpret topic models, in: Neural Information Processing Systems, 2009.
- [35] D. Newman, J. H. Lau, K. Grieser, T. Baldwin, Automatic evaluation of topic coherence, in: Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics, 2010.

- [36] D. M. Blei, Probabilistic topic models, *Communications of the ACM* 55 (4) (2012) 77–84.
- [37] W. D. Gray, R. M. Young, S. S. Kirschenbaum, Introduction to this special issue on cognitive architectures and human-computer interaction, *Human-Computer Interaction* 12 (4) (1997) 301–309.
- [38] Z. Duric, W. Gray, R. Heishman, F. Li, A. Rosenfeld, M. Schoelles, C. Schunn, H. Wechsler, Integrating perceptual and cognitive modeling for adaptive and intelligent human-computer interaction, *Proceedings of the IEEE* 90 (7) (2002) 1272–1289.
- [39] S. Lallée, V. Vouloutsi, S. Wierenga, U. Pattacini, P. Verschure, EFAA: a companion emerges from integrating a layered cognitive architecture, in: *ACM/IEEE International Conference on Human-robot Interaction*, 2014.
- [40] P. E. Baxter, J. de Greeff, T. Belpaeme, Cognitive architecture for human–robot interaction: towards behavioural alignment, *Biologically Inspired Cognitive Architectures* 6 (2013) 30–39.
- [41] B. Yang, L. Zhou, Z. Deng, C-HMAX: Artificial cognitive model inspired by the color vision mechanism of the human brain, *Tsinghua Science and Technology* 18 (1) (2013) 51–56.
- [42] M. Halbrügge, ACT-CV: Bridging the gap between cognitive models and the outer world, *Grundlagen und anwendungen der mensch-maschine-interaktion* (2013) 205–210.
- [43] H.-H. Nagel, Steps toward a cognitive vision system, *AI Magazine* 25 (2) (2004) 31–50.
- [44] A. Lawitzky, D. Althoff, C. F. Passenberg, G. Tanzmeister, D. Wollherr, M. Buss, Interactive scene prediction for automotive applications, in: *Intelligent Vehicles Symposium*, 2013, pp. 1028–1033.

- [45] J. L. Crowley, Things that see: Context-aware multi-modal interaction., in: Cognitive Vision Systems, 2006, pp. 183–198.
- [46] O. Brdiczka, J. L. Crowley, P. Reignier, Learning situation models in a smart home, IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics 39 (1) (2009) 56–63.
- [47] F. Xu, T. L. Griffiths, Probabilistic models of cognitive development: Towards a rational constructivist approach to the study of learning and development, Cognition 120 (2011) 299–301.
- [48] N. Bencomo, A. Belaggoun, V. Issarny, Dynamic decision networks for decision-making in self-adaptive systems: A case study, in: International Symposium on Software Engineering for Adaptive and Self-Managing Systems, 2013.
- [49] M. Kafai, B. Bhanu, Dynamic Bayesian networks for vehicle classification in video, IEEE Transactions on Industrial Informatics 8 (1) (2012) 100–109.
- [50] N. Chater, J. B. Tenenbaum, A. Yuille, Probabilistic models of cognition: where next?, Trends in Cognitive Sciences 10 (7) (2006) 292–293.
- [51] D. M. Blei, A. Y. Ng, M. I. Jordan, Latent dirichlet allocation, Journal of Machine Learning Research 3 (2003) 993–1022.
- [52] T. L. Griffiths, M. Steyvers, Finding scientific topics, in: National Academy of Sciences, 2004.
- [53] I. Porteous, D. Newman, A. Ihler, A. Asuncion, P. Smyth, M. Welling, Fast collapsed Gibbs sampling for latent dirichlet allocation, in: ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2008.
- [54] C. C. Musat, J. Velcin, S. Trausan-Matu, M.-A. Rizoio, Improving topic evaluation using conceptual knowledge, in: International Joint Conference on Artificial Intelligence, 2011.
- [55] D. Blei, J. Lafferty, Correlated topic models, in: Neural Information Processing Systems, 2006.

- [56] M. Dondar, F. Akova, A. Qi, B. Rajwa, Bayesian nonexhaustive learning for on-line discovery and modeling of emerging classes, in: International Conference on Machine Learning, 2012.
- [57] L. Gorelick, M. Blank, E. Shechtman, M. Irani, R. Basri, Actions as space-time shapes, IEEE Transactions on Pattern Analysis and Machine Intelligence 29 (12) (2007) 2247–2253.
- [58] I. Laptev, On space-time interest points, International Journal of Computer Vision 64 (2-3) (2005) 107–123.